

Should our cognitive biases have moral weight?

In [a classic piece of psychology](#), Kahneman and Tversky ask people what to do about a fatal disease that 600 people have caught. One group is asked whether they would administer a treatment that would definitely save 200 people's lives or one with a 33% chance of saving 600 people. The other group is asked whether they would administer a treatment under which 400 people would definitely die or one where there's a 33% chance that no one will die.

The two questions are the same: saving 600 people means no one will die, saving just 200 means the other 400 will die. But people's responses were radically different. The vast majority of people chose to save 200 people for sure. But an equally large majority chose to take the chance that no one will die. In other words, just changing how you describe the option — saying that it saves lives rather than saying it leaves people to die — changes which option most people will pick.

In the same way that Festinger, et. al. showed that our intuitions are biased [by our social situation](#), Kahneman and Tversky demonstrated that humans suffer from consistent cognitive biases as well. In a whole host of examples, they showed people behaving in a way we wouldn't hesitate to think was irrational — like changing their position on whether to administer a treatment based on what it was called. (I think a similar problem affects our intuitions about killing versus letting die.)

This is a major problem for people like Frances Kamm, who think our moral philosophy must rely on our intuitions. If people consistently and repeatedly treat things differently based on what they're called, are we forced to give that moral weight? Is it OK to administer a treatment when it's described as saving people, but not when it's described as not saving enough? Surely moral rules should meet some minimal standard of rationality.

This problem affects a question close to Kamm's work: what she calls the Problem of Distance in Morality (PDM). Kamm says that her intuition consistently finds that moral obligations attach to things that are close to us, but not to things that are far away. According to her, if we see a child drowning in a pond and there's a machine nearby which, for a dollar, will scoop him out, we're morally obligated to give the machine a dollar. But if the machine is here but the scoop and child are on the other side of the globe, we don't have to put a dollar in the machine.

But, just as with how things are called, our intuitions about distance suffer from cognitive biases. Numerous studies have shown that the way we think about things nearby is radically different from the way we think about things far away. In [one study](#), Indiana University students did better on a creativity test when they were told the test was devised by IU students studying in Greece than when they were told it was devised by IU students studying in Indiana.

It's a silly example, but it makes the point. If our creativity depends on whether someone mentions Greece or Purdue, it's no surprise our answers to moral dilemmas depend on whether they take place in the US or China. But surely these differences have no more moral validity than the ones that result from Tversky's experiment — they're just an unfortunate quirk of how we're wired. Rational reflection — not faulty intuitions — should be the test of a moral theory.

You should follow me on twitter [here](#).

January 8, 2010